arXiv:1210.2376v1 [physics.soc-ph] 8 Oct 2012

# Interdependence and Predictability of Human Mobility and Social Interactions

Manlio De Domenico, Antonio Lima, Mirco Musolesi

*School of Computer Science, University of Birmingham*
*Edgbaston B15 2TT, Birmingham, United Kingdom*

**Abstract**

Previous studies have shown that human movement is predictable to a certain extent at different geographic scales. Existing prediction techniques exploit only the past history of the person taken into consideration as input of the predictors.

In this paper, we show that by means of multivariate nonlinear time series prediction techniques it is possible to increase the accuracy of movement forecasting by considering movements of friends or people with correlated mobility patterns (i.e., characterised by high mutual information) as input of the predictor. Finally, we evaluate the proposed techniques on the Nokia Mobile Data Challenge and Cabspotting datasets.

*Keywords:* mobility prediction, mutual information, nonlinear timeseries analysis

## 1. Introduction

The study of the interdependence of human movement and social ties of individuals is one of the most interesting research areas in computational social science [1]. Previous studies have shown that human movement is predictable to a certain extent at different geographic scales [2, 3, 4]. The potential uses of these prediction techniques are various, including practical ones, such as content dissemination of location-aware information, e.g., tar-

*Email addresses:* m.dedomenico@cs.bham.ac.uk (Manlio De Domenico), a.lima@cs.bham.ac.uk (Antonio Lima), m.musolesi@cs.bham.ac.uk (Mirco Musolesi)

geted advertisements in sponsored mobile applications or in search results performed from mobile phones [5].

In this paper we show how it is possible to improve mobility prediction by exploiting the correlation between movements of friends and acquaintances. This can be seen as a process of discovering correlation patterns in networks of linked social information and geographic data. It is possible to exploit such correlations for prediction and inference of aspects related to user behaviour, namely their movements and their social interactions (either physical and distant). In particular, in our analysis we exploit and adapt the concept of *mutual information* [6] in order to quantify correlation and provide a *practical* method for the selection of additional data to improve the accuracy of movement forecasting. We also show how this quantity correlates to different types of social interactions of friends and acquaintances. This paper extends the findings presented in our submission [7] to the Nokia Mobile Data Challenge competition [8].

More specifically, the contributions of this work can be summarised as follows:

- We first show that by means of a multivariate nonlinear predictor [9] we are able to achieve a very high degree of accuracy in forecasting future user geographic locations in terms of longitude and latitude. We compare it with traditional linear prediction techniques (such as ARMA [10]) and we show that these are not able to capture the dynamics of individuals in the geographic space.

- We discuss how the concept of mutual information can be used to quantify the correlation between two mobility traces and we demonstrate that it is possible to exploit movement data of friends and acquaintances, where such information is available.

- Finally, we study how the correlation measured through mutual information of mobility traces of two individuals, can be used to improve human prediction movement dramatically, also discussing the correlation between human mobility and social ties.

The key findings of our analysis are the following: 1) mobility correlation can be used to improve movement forecasting by exploiting mobility data of friends; 2) correlated movement is linked to the existence of physical or distant social interactions and vice versa.

We evaluate these findings on two datasets. The first dataset, which was provided for the Nokia Mobility Data Challenge (NMDC), contains information related to 39 users [8], including the following: GPS traces, telephone numbers, call and SMS history, Bluetooth and WLAN history. We use the information of 25 of them, since the dataset does not include phone numbers for 14 of them; therefore, it is not possible to detect if and when phone calls occur between them. We use GPS traces to analyse the movement of the users.

The second dataset we analyse is Cabspotting [11], containing mobility traces of about 500 taxis driving around San Francisco for 30 days. We restrict our analysis to the 178 taxis with mobility traces longer than 25000 GPS readings. For this dataset we have no information about relations between taxi drivers (such as friendship connections or co-affiliation).

The paper is organised as follows. In Section 2 we firstly introduce multivariate nonlinear time series prediction techniques and their application to our dataset. Then, in Section 3 we discuss how mutual information can be used to measure the correlation between the movement of two users. Section 4 focusses on the analysis of the performance improvement that is possible to obtain by considering the traces of highly correlated users as inputs of the predictors. In Section 5, we discuss our findings also outlining some future directions. Section 6 concludes the paper, summarising our contributions.

## 2. Multivariate Nonlinear Time Series Prediction

We now present how we apply nonlinear time series prediction methods to the problem of forecasting the future GPS coordinates of the users, given the past movement history as an input. We will then extend this model by considering also the movement of other users (in particular friends, in the case of the NMDC dataset) as input of the nonlinear predictor.

### *2.1. Overview*

We model the position of a user on the Earth with a time-varying four-dimensional state vector $\mathbf{s}_n$ with the following dimensions: hour of the day $h_n$, latitude $\phi_n$, longitude $\lambda_n$ and altitude[1] $\xi_n$. The prediction of the future states of vector $\mathbf{s}_n$ can be performed using different approaches [9]. We

[1]The corresponding time series is available only in the NMDC dataset: in the case of Cabspotting data we use a time-varying three-dimensional state vector.

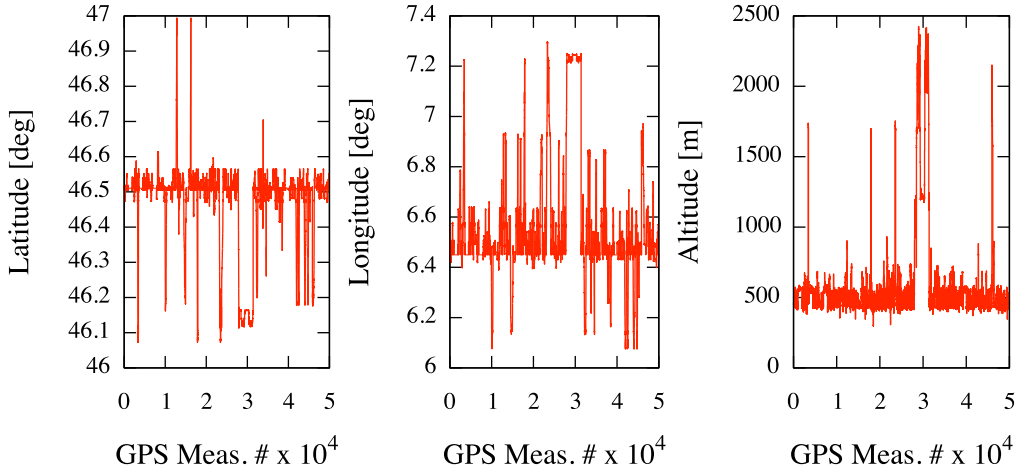


Figure 1: Time series from the Nokia Mobility Data Challenge dataset, corresponding to the movements of user 179. No periodic behaviour is apparent in the movement traces of the user.

choose the method based on the reconstruction of the phase space of $\mathbf{s}_n$ by means of the delay embedding theorem, since this is considered the best state-of-the-art solution to this problem. While the scalar sequence of coordinates may appear completely non deterministic, it is possible to uncover the characteristics of its dynamic evolution by analysing sub-sequences of the time series itself. In order to investigate the structure of the original system, the time series values must be transformed in a sequence of vectors with a technique called delay embedding. For a univariate time series measurement $x_n$ of a $d-$dimensional dynamical system, the Takens' embedding theorem [12] allows to reconstruct a $m-$dimensional space ($m \geq 2d+1$) with the same dynamical characteristics of the original phase space. The key idea is to build a delay vector $\mathbf{x}_n$ by using delayed measurement defined as follows:

$$\mathbf{x}_n \equiv (x_{n-(m-1)\tau}, x_{n-(m-2)\tau}, ..., x_{n-\tau}, x_n), \tag{1}$$

where $\tau$ is a time delay. Hence, the reconstruction depends on the two parameters $m$ and $\tau$, which have to be estimated. This technique can be extended to the case of the embedding of a multivariate time series[2] [14].

[2]We refer to [13] (and references therein) for an overview of practical applications of multivariate embedding.

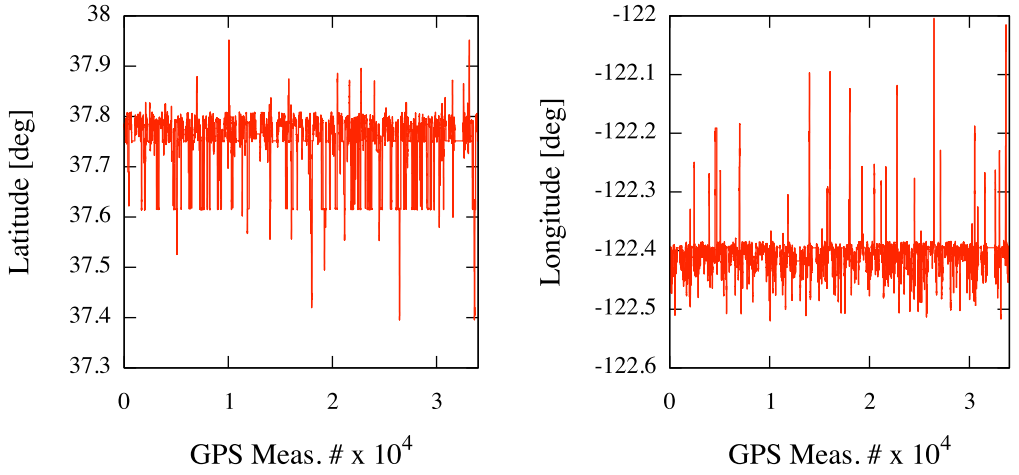


Figure 2: Time series from the Cabspotting dataset, corresponding to the movements of taxi *abgibo*.

Under the hypotheses of Takens' theorem, i.e., non-noisy time series of infinite length, the underlying dynamics can be fully reconstructed by using only univariate measurements of the dynamical system of interest. Unfortunately, real-world measurement are noisy and with finite length: hence, the phase space reconstruction is more precise if multivariate measurements of the dynamical system under investigation are performed.

Let us indicate with $N$ the number of measurements corresponding to an $M-$dimensional time series $\mathbf{y}_1$, $\mathbf{y}_2$, ..., $\mathbf{y}_N$, with $\mathbf{y}_i \equiv (y_{1,i}, y_{2,i}, ..., y_{M,i})$ and $i = 1, 2, ..., N$. The resulting delay vector is

$$\begin{aligned} \mathbf{v}_n \equiv & (y_{1,n-(m_1-1)\tau_1}, y_{1,n-(m_1-2)\tau_1}, ..., y_{1,n}, \\ & y_{2,n-(m_2-1)\tau_2}, y_{2,n-(m_2-2)\tau_2}, ..., y_{2,n}, \\ & ... \\ & y_{M,n-(m_M-1)\tau_M}, y_{1,n-(m_M-2)\tau_M}, ..., y_{M,n}), \end{aligned} \tag{2}$$

where $m_j$ and $\tau_j$, $j = 1, 2, ..., M$ are respectively the embedding and time delays corresponding to each component of the multivariate time series.

Intuitively, this method searches the past history to find and extract sequences of values that are very similar to the recent history. Assuming a certain degree of determinism in the system, the assumption is that, given a

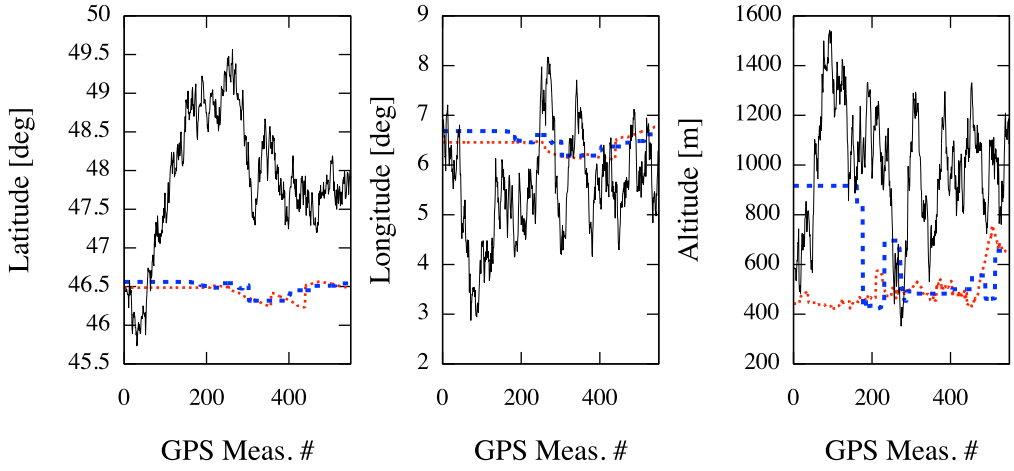


Figure 3: Multivariate nonlinear prediction of user 179 mobility in the geographic space: the first 600 predictions, corresponding to about 60 hours, are shown. The dotted line represents the true movements, the solid line indicates the prediction within an ARMA model, while the dashed line indicates the data obtained by means of a multivariate nonlinear predictor.

certain state (in our case geographic coordinates), there is a strong probability that this will be followed by the same next state.

### *2.2. Evaluation*

#### *2.2.1. Linearity Analysis*

The complexity of the time series taken into consideration in our study is apparent by observing the two representative examples shown in Fig. 1 and Fig. 2. The figures show thousands of time-ordered GPS measurements corresponding to the position on the Earth of user 179 (NMDC dataset) and taxi *abgibo* (Cabspotting dataset), respectively.

We firstly apply linear prediction models to these time series. The time series appear rather noisy with alternating spikes, nearly flat values, corresponding to stationary points, and fluctuation around an average value. We try to model such movements in the space with a simple multivariate AR($p$) + noise process.

As for the order $p$ of the multivariate autoregressive model that best approximates the original time series, we choose the one that minimises an information criterion, according to Akaike [15] and Schwarz [16]. We find

that $p = 24$ provides the best approximation. Hence, we use such a model to perform a multivariate linear forecasting of 1000 GPS measurements for user 179 (NMDC dataset). We validate the model by comparing the latest 1000 real GPS measurements against the forecasted ones[3]. The results are shown in Fig. 3, where the real movements are indicated with dots and the forecasting with the linear model is indicated by the solid line. It is evident that the forecasting is not in agreement with observations. In fact, the prediction error on the position (latitude and longitude) is of the order of 3°, whereas the error on the altitude is generally larger than 600 m.

However, although the time series are not regularly sampled, we find that they show some features typical of deterministic dynamics contaminated by noise. In fact, preliminary inspection of phase space reconstruction by means of Takens' embedding theorem shows an underlying structure, typical of deterministic dynamical systems. This aspect will be addressed more quantitatively in the remainder of the paper.

*2.2.2. Estimation of the Embedding Dimension and Time Delay*

Although several methods have been proposed to estimate the values of embedding and time delay, in our analysis we consider the same time delay $\tau$ for all the series. In fact, for a given user, we have found that the time delay $\tau_{min}$ corresponding to the first local minimum of the average mutual information [17], generally adopted to estimate $\tau$ in the univariate case, is of the same order of magnitude for any component. As a representative example, in Fig. 4 we show the distribution of $\tau_{min}$ obtained from the time series of latitude and longitude of taxis in the Cabspotting dataset.

This fact has also practical implications, since it simplifies the application of this methodology for the analysis of our data. The optimal embedding dimension is estimated by exploiting the method of false nearest neighbours [18, 9, 19] in the case of multivariate embedding [20]. For any point in the data, an $m^\star$-dimensional phase space is considered and the number of false nearest neighbours, i.e., points which are neighbours in the $m^\star-$dimensional space but not in the $(m^\star + 1)-$dimensional one, is estimated. The desirable embedding dimension $m$ is such that the percentage of false nearest neighbours is small, e.g., below 5%. Any efficient algorithm for counting nearest

[3]The latest 1000 real GPS measurements have not been included in the procedure adopted to estimate the best order $p$.

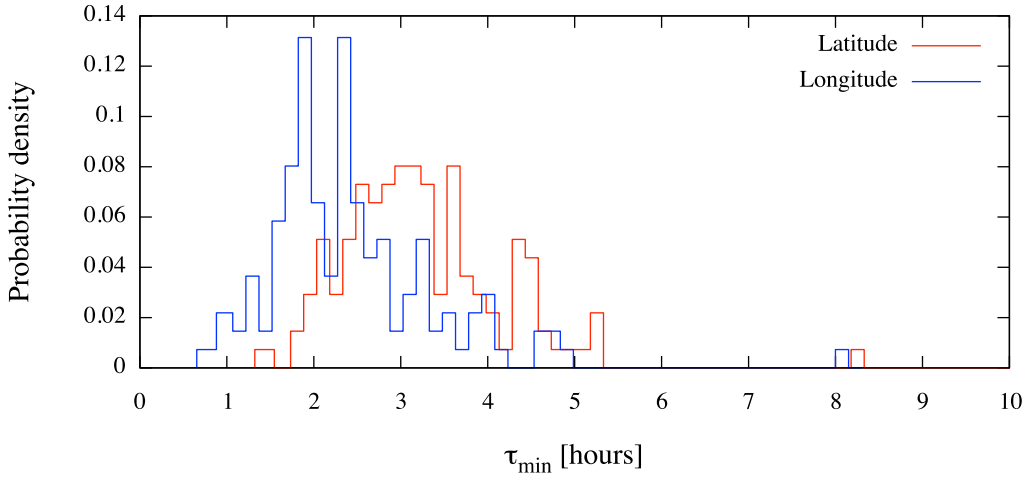


Figure 4: *Cabspotting dataset.* Distribution of the time delay $\tau_{min}$ corresponding to the first local minimum of the average mutual information obtained from the time series of latitude and longitude. The average value of $\tau_{min}$ is approximately 3 hours.

neighbours is allowed: in particular, we adopt the method implemented in the TISEAN software [21]. In the left panel of Fig. 5 we show the fraction of false nearest neighbours as a function of $m$, obtained from mobility traces in the Cabspotting dataset. For any trace, the optimal embedding dimension is close to 30, confirming that the underlying dynamics is very similar[4]. The false nearest neighbour method alone is not able to distinguish between deterministic and stochastic processes on an absolute level [19]: however, it is among the state-of-the-art solutions [22, 23, 24, 25, 26] that can be reliably used to asses the nonlinearity of time series by means of a statistical test with surrogate data.

*2.2.3. Analysis of Multivariate Surrogates*

Given a multivariate time series $\{\mathbf{s}_n\}$, we produce a set $\{\hat{\mathbf{s}}_n^{(i)}\}$, $i = 1, 2, ..., N$, of $N$ multivariate surrogates of $\{\mathbf{s}_n\}$. The surrogates are synthetic time series, built from $\{\mathbf{s}_n\}$, preserving both statistical and linear features of the original time series, as probability distribution and autocorrelation, while removing the effects of nonlinearities and nonstationarities, if any. In partic-

[4]We find only a few exceptions whose number represents less than 5% of the mobility traces in the whole dataset.

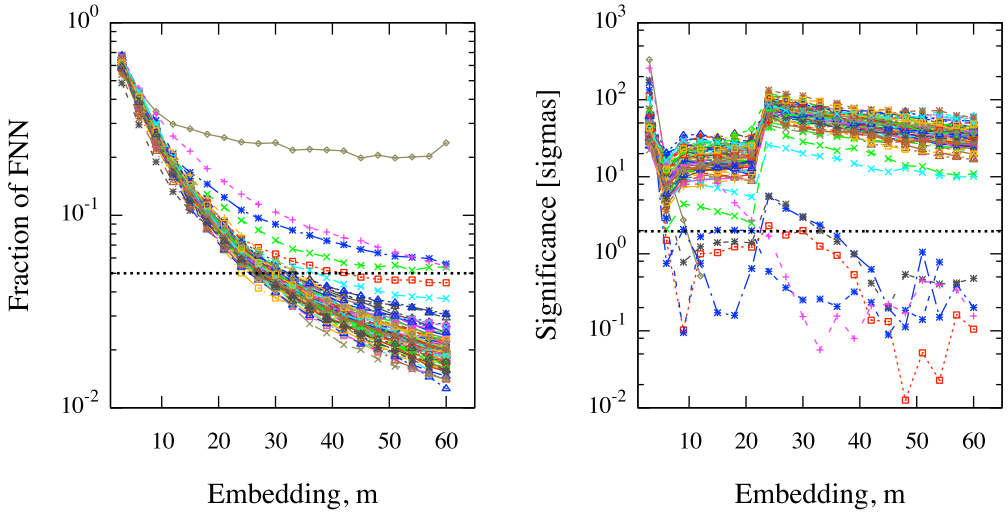


Figure 5: *Cabspotting dataset.* Left panel: fraction of false nearest neighbours as a function of embedding dimension $m$; each curve corresponds to a different mobility trace and the threshold chosen to estimate the optimal embedding is indicated by the dotted horizontal line. Right panel: significance in rejecting the null hypothesis by using the fraction of false nearest neighbours as discriminator in the surrogate test (see the text for further details); each curve corresponds to a different mobility trace and the size of the statistical test is indicated by the dotted horizontal line.

ular, we adopt the iterative amplitude-adjusting Fourier transform (IAAFT) scheme [27, 28] to build surrogates. Hence, we choose the fraction of false nearest neighbours as discriminator to test the null hypothesis that the mobility traces can be described by a linear stochastic model. Let us indicate with $q(m)$ the value of the discriminator obtained for an embedding dimension $m$ from the observed multivariate time series, and with $\hat{q}^{(i)}(m)$ the values of the discriminator obtained from surrogates. Our numerical experiments indicate that the distribution of $\hat{q}^{(i)}(m)$ is described with a reasonable approximation by a Gaussian function with average $\mu_{\hat{q}}(m)$ and variance $\sigma^2_{\hat{q}}(m)$. This fact allows us to define the quantity

$$\Sigma(m) \equiv \frac{|q(m) - \mu_{\hat{q}}(m)|}{\sigma_{\hat{q}}(m)} \tag{3}$$

as a measure of *significance*. In this case, if the null hypothesis is true then the $p-$value of observing a significance equal or larger than $\Sigma(m)$ is given by $p(m) = \mathrm{erfc}[\Sigma(m)/\sqrt{2})]$. We fix *a priori* the size $\alpha = 0.05$ of our hypothesis testing: if $p(m) < \alpha$ (or, equivalently, if $\Sigma(m) > 1.96$) the null

hypothesis that mobility can be described by a linear stochastic model is rejected with 95% confidence level (CL). In the right panel of Fig. 5 we show the significance as a function of $m$ for mobility traces in the Cabspotting dataset. Remarkably, the significance is much larger than 1.96 for all traces despite a few exceptions, independently from the embedding dimension chosen for the reconstruction. Hence, we can conclude that human mobility exhibits a strong nonlinear dynamics. Moreover, the existence of short-term correlations, as indicated by the average mutual information analysis, and of decreasing fraction of false nearest neighbours for increasing embedding dimension suggests that such a dynamics should have a deterministic component potentially contaminated by a stochastic dynamics.

*2.2.4. Analysis of Prediction Errors*

Dealing with nonlinear dynamical systems with a potential deterministic component, we adopt a method which exploits such features to predict the future movements of users in the NMDC dataset and of taxis in the Cabspotting dataset. The multivariate nonlinear prediction (MNP) is performed by approximating the dynamics locally in the phase space by a constant (see [29] for further information). In the delay embedding space, all the points in the neighbourhood $\mathcal{U}_n$ of the state $\mathbf{v}_n$ are taken into account in order to predict the coordinates at time $n+k$. Hence, the forecast $\hat{\mathbf{v}}_{n+k}$ for $\mathbf{v}_{n+k}$ is given by

$$\hat{\mathbf{v}}_{n+k} = \frac{1}{|\mathcal{U}_n|} \sum_{\mathbf{v}_j \in \mathcal{U}_n} \mathbf{v}_{j+k}, \tag{4}$$

i.e., the average over the states which correspond to measurements $k$ steps ahead of the neighbours $\mathbf{v}_j$.

Hence, we use MNP to forecast the same 1000 GPS measurements previously discussed in the case of NMDC dataset. Again, we validate the model by comparing the latest 1000 real GPS measurements against the forecasted ones. The results for user 179 are shown in Fig. 3, where the real movements are indicated with triangles and the forecasting with the nonlinear method is indicated by the dashed line. The number of nearest neighbours used to build the neighbourhood $\mathcal{U}_n$ has been kept fixed to 10. Intriguingly, the nonlinear forecasting is in excellent agreement with observations of latitude and longitude, with a global position prediction error equal to 0.19°, and in good agreement with the altitude coordinate, with a global altitude forecasting error equal to 219.43 m.

The global error on the time series prediction is estimated separately for each component using the following formula:

$$e_j = \sqrt{\frac{1}{N}\sum_{n=1}^{N}(\hat{s}_{j,n} - s_{j,n})^2}, \tag{5}$$

with $j = 1, 2, ..., M$ with $M = 4$, $N = 1000$. The overall error between the predicted position and the real one is given by the geodesic distance.

## 3. Mutual Information and Movement Correlation

In this section, we will briefly introduce the concept of mutual information and we will show how this quantity can be exploited in our analysis to measure the correlation between the movement of different individuals. In the following section, we will then discuss how mutual information can be used to select mobility data of other users that can be exploited as inputs of the nonlinear predictors in order to improve the prediction accuracy.

### *3.1. Overview*

Let us assume that $\mathbf{X}$ and $\mathbf{Y}$ are two multivariate stochastic variables, and let us indicate with $P_{\mathbf{X}}(\mathbf{x})$ and $P_{\mathbf{Y}}(\mathbf{y})$, respectively, the corresponding Probability Density Functions (PDF). The joint probability is indicated by $P_{\mathbf{XY}}(\mathbf{x}, \mathbf{y})$. The mutual information $\mathcal{I}(\mathbf{X}, \mathbf{Y})$ between such two variables is defined as follows:

$$\mathcal{I}(\mathbf{X}, \mathbf{Y}) = \sum_{\mathbf{x}\in\mathbf{X}}\sum_{\mathbf{y}\in\mathbf{Y}} P_{\mathbf{XY}}(\mathbf{x}, \mathbf{y}) \log \frac{P_{\mathbf{XY}}(\mathbf{x}, \mathbf{y})}{P_{\mathbf{X}}(\mathbf{x})P_{\mathbf{Y}}(\mathbf{y})}. \tag{6}$$

The mutual information[5] quantifies how much information the variable $\mathbf{Y}$ provides about the variable $\mathbf{X}$. For this reason, it can be used as an estimator of the amount of correlation between $\mathbf{X}$ and $\mathbf{Y}$. In fact, if the two variables are totally uncorrelated then $P_{\mathbf{XY}}(\mathbf{x}, \mathbf{y}) = P_{\mathbf{X}}(\mathbf{x})P_{\mathbf{Y}}(\mathbf{y})$ and $\mathcal{I}(\mathbf{X}, \mathbf{Y}) = 0$.

In our analysis $\mathbf{X}$ represents the motion of a user on the Earth, the random samples $\mathbf{x}$ drawn from $\mathbf{X}$ correspond to geographic coordinates, whereas the PDF of $\mathbf{x}$ quantifies the fraction of time spent by the user in a particular position.

[5]The units of mutual information are nats when the natural logarithm is used.

We use the mutual information to quantify the amount of correlation between the motion of different users, or, equivalently, how much information the motion $\mathbf{Y}$ provides about the motion $\mathbf{X}$.

*3.2. Evaluation*

In the NMDC dataset, we say that two individuals are friends or acquaintances if one of them is in the other's address book. In Fig. 6 the two-dimensional PDF of positions occupied by four different users is shown. Users 063 and 123 are friends or acquaintances, while users 026 and 127 are not.

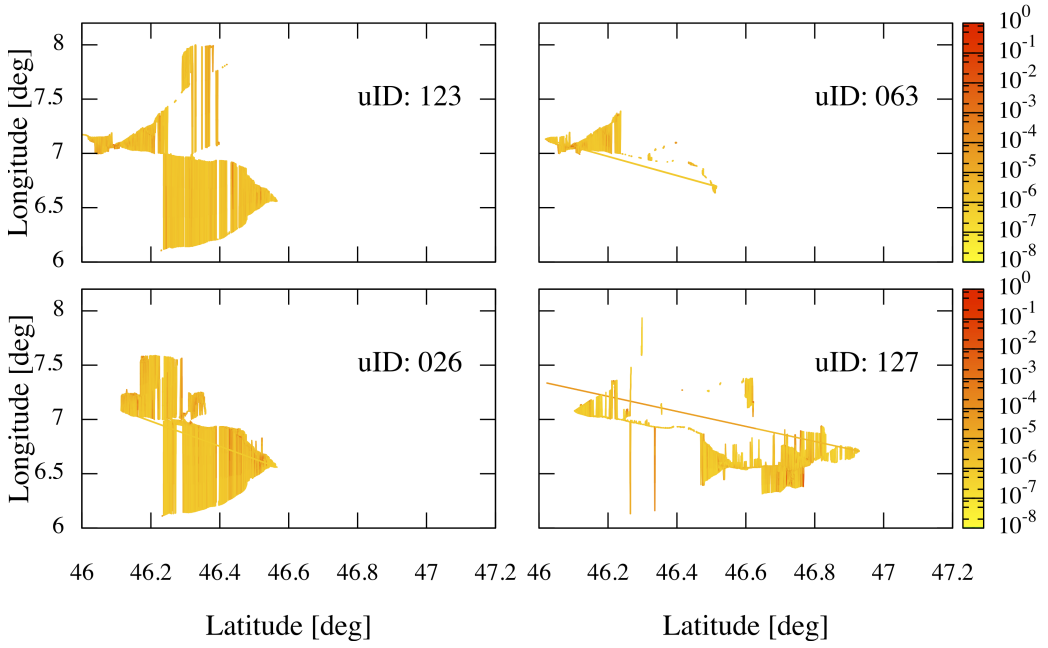


Figure 6: *NMDC dataset.* PDF of locations occupied by four different users. Top: users are friends or acquaintances. We say that two individuals are friends or acquaintances if one of them is in the other's address book. Bottom: users are not friends or acquaintances. The colour indicates the frequency of occupation.

It is worth noting that friendship or acquaintanceship are only sufficient conditions to have a high mutual information, but they are not necessary. In fact, the mutual information between two users who are not friends or acquaintances can be still high if the two users behave similarly in space and time, i.e., if their motion is similar for some reasons (e.g., a pair of

students of the same university department or a pair of colleagues of the same company).

## 4. Exploiting Movement Correlation and Social Ties to Improve Prediction Accuracy

We now discuss how mobility traces of individuals that have correlated geographic patterns and social ties can be used to improve the accuracy of movement forecasting.

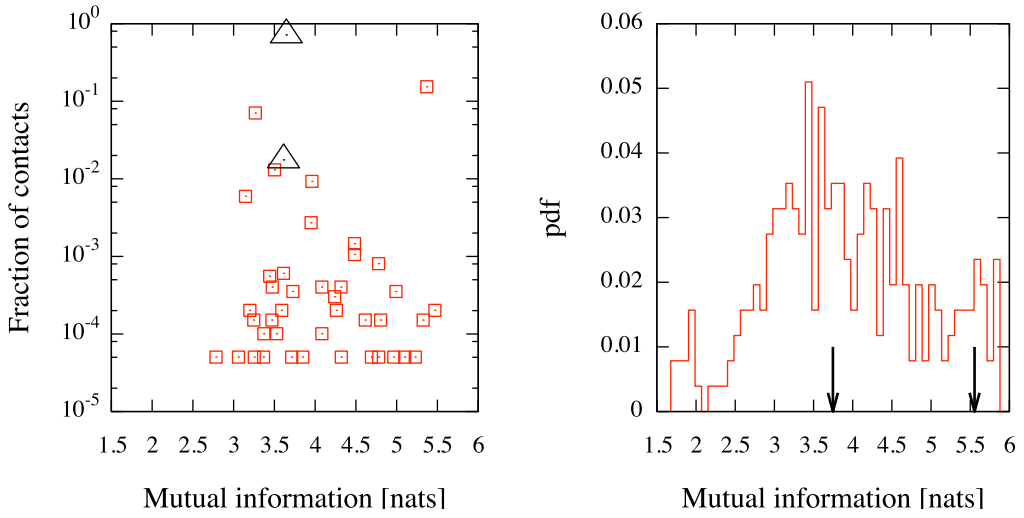


Figure 7: *NMDC dataset.* Left panel: scatter plot of the fraction of contacts vs the mutual information estimated for pairs of users with at least one contact, where triangles indicate the two pairs of users connected by social ties in the dataset. Right panel: pdf of mutual information estimated for pairs of users with no contacts at all, where arrows indicate the value of mutual information for the only two pairs of user with social ties, and no contacts, in the dataset.

### *4.1. Our Approach at a Glance*

Our approach can be summarised as follows: assuming that we want to predict the movement of person/entity $A$, instead of having only the vector describing the location of $A$ as input, we will also consider the movement history of another person/entity $B$, characterised by mobility patterns that are strongly correlated to those of the user we would like to predict. This measure is given by the mutual information introduced in the previous section.

From a mathematical point of view, the idea is to use a 8-dimensional vector that is given by the juxtaposition of the two time-varying state vectors representing the states (time-stamped GPS coordinates) of $A$ and $B$, which we indicate with $\mathbf{s}_{n_A}$ and $\mathbf{s}_{n_B}$, as inputs of the multivariate nonlinear predictor.

*4.2. Evaluation*

In both datasets we find that by using additional traces of pairs with high correlation, the accuracy of the prediction improves consistently. In the case of NMDC, the improvement is of at least one order of magnitude (and often of two orders of magnitude) with respect to the prediction based on only single traces. Moreover, it is interesting to note that social ties can also be used to select the user for the additional traces as input. In fact, we find that if we select mobility patterns of individuals that are in the address book of the user, the performance of the predictor improves dramatically. At the same time, we would like to stress the fact that the NMDC dataset contains a small number of users, therefore it is difficult to make claims about the general validity of these findings. However, we find the same results for much larger Cabspotting dataset. In this dataset, it is not possible to use social ties [6], but we find that if we select mobility patterns of taxis whose mutual information is high, the performance of the predictor improve drastically.

Hence, we perform the same analysis described in Section 2, but including the time series of movements corresponding to other users in the multivariate nonlinear prediction. The global prediction error, defined by Eq. (5), of position and altitude is reported in Tab. 1 for the three cases of study in the NMDC dataset. As shown in this table, we observe that the additional information provided by the movement of a user socially linked to that taken into consideration improves the prediction by more than one order of magnitude with respect to the case of users who are not socially linked to each other.

For each pair of users in the NMDC dataset, we count the total number of Bluetooth contacts and calls, then we estimate their mutual information defined by Eq. (6). In order to quantify the amount of correlation between the fraction of contacts and the mutual information, we build a scatter plot between these two observables. The result is shown in the left panel of Fig. 7,

[6] In theory, it might be interesting to investigate the influence of the social ties between taxi drivers, but this information unfortunately is not available in the dataset.

| Nodes | Social link | Position Error | Altitude Error |
|---|---|---|---|
| 026 127 | None | 0.167° | 66.33 m |
| 063 123 | Present | 0.011° | 20.95 m |
| 094 009 | Present | 0.003° | 5.57 m |

Table 1: *NMDC dataset.* Global error, defined by Eq. (5), on the prediction of position and altitude for pairs of users connected through social links (defined as presence in the address book of the user).

by considering only pair of users with at least one contact. The points corresponding to pairs of users with social ties are also shown (triangles). In the right panel of Fig. 7, we show the PDF of mutual information obtained by considering only pairs of users with no contacts at all. The mutual information corresponding to pairs of users with social ties is shown (arrows). Even if these plots show interesting correlations for this specific dataset, we believe no generalisations can be drawn from them, because of the lack of sufficient statistics.

Hence, we perform the same analysis by exploiting the mobility patterns of taxis in the Cabspotting dataset, which contains a larger statistical sample. In this case, the global prediction error, defined by Eq. (5), refers only to latitude and longitude. Moreover, we investigate the evolution of the global prediction error by estimating how it changes versus time. More specifically, we define the time-varying global prediction error for each component as

$$e_j(t) = \sqrt{\frac{1}{t}\sum_{n=1}^{t}(\hat{s}_{j,n} - s_{j,n})^2}, \tag{7}$$

with $j = 1, 2$ and $t$ indicating the prediction interval. Hence, the overall error between the predicted position and the real one at time $t$ is given by $e(t) = \sqrt{e_1^2(t) + e_2^2(t)}$. In order to investigate the quality of our prediction, we study the ratio of $e(t)$ with respect to the global statistical uncertainty $\sigma_{data}$ on the position of the taxi. In fact, as long as the ratio $e(t)/\sigma_{data}$ is equal or smaller than one, or, equivalently, if $e(t) \leq \sigma_{data}$, the prediction at time $t$ is within the statistical uncertainty and, therefore, the performance of our predictor can be considered satisfactory. In Fig. 8 we show the cumulative distribution of the values of the ratio obtained from mobility traces in the Cabspotting dataset. In particular, we show the distributions corresponding

to the predicted positions after 5 minutes and 30 minutes. The three curves correspond to prediction involving: a) only the past history of each single taxi ("Single"), b) the history of any pairs of taxi whose mobility patterns show a low mutual information ("Combined, Low MI") and c) the history of any pairs of taxi whose mobility patterns show a high mutual information ("Combined, High MI"). It is worth remarking that the mutual information is not an upper-bounded measure of correlation: hence, we define "high MI" all pairs of mobility patterns whose mutual information is distributed among the highest 5% of values, and "low MI" the remaining pairs of mobility patterns. In both panels of Fig. 8, we can observe that the prediction improves when combining pairs of correlated mobility patterns. Moreover, it is intriguing that our method is able to predict in the 80% of cases the movements of taxis for the next 30 minutes, with an error equal or smaller than the statistical uncertainty of their mobility patterns.

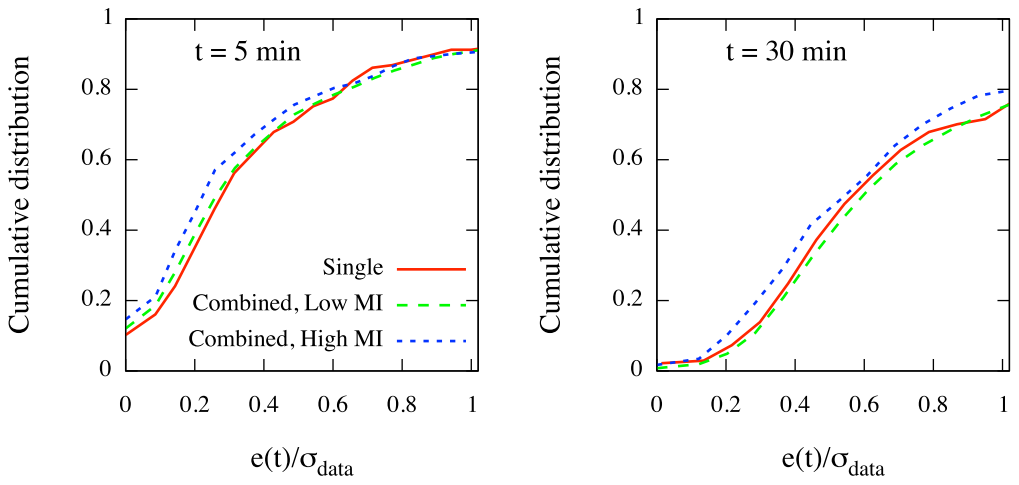


Figure 8: *Cabspotting dataset.* Cumulative distribution of the values of the ratio between the global prediction error $e(t)$ and the statistical uncertainty of the mobility traces $\sigma_{data}$. We show the distributions corresponding to the predicted positions after 5 minutes (left panel) and 30 minutes (right panel). The three curves correspond to prediction involving: a) only the past history of each single taxi ("Single"), b) the history of any pairs of taxi whose mobility patterns show a low mutual information ("Combined, Low MI") and c) the history of any pairs of taxi whose mobility patterns show a high mutual information ("Combined, High MI"). .

Since the prediction is acceptable when the ratio $(e(t)/\sigma_{data})$ is below one, we investigate how the fraction of mobility traces satisfying this requirement,

i.e., $P(e(t)/\sigma_{data} < 1)$, changes over time. In Fig. 9 we show this temporal evolution, with 90% confidence bands around the average values. Due to the lack of statistics (137 mobility traces) in the "Single" traces prediction, the bands are wider than for other cases (9316 mobility traces). All curves show decreasing behaviour for increasing prediction interval, as expected. In fact, the "combined, high MI" predictor performs equal or better than others up to about 90 minutes. It is worth mentioning that the forecasting of every predictor is within the statistical uncertainity ($e(t) \leq \sigma_{data}$) for more than 50% of mobility traces considered up to 3 hours.

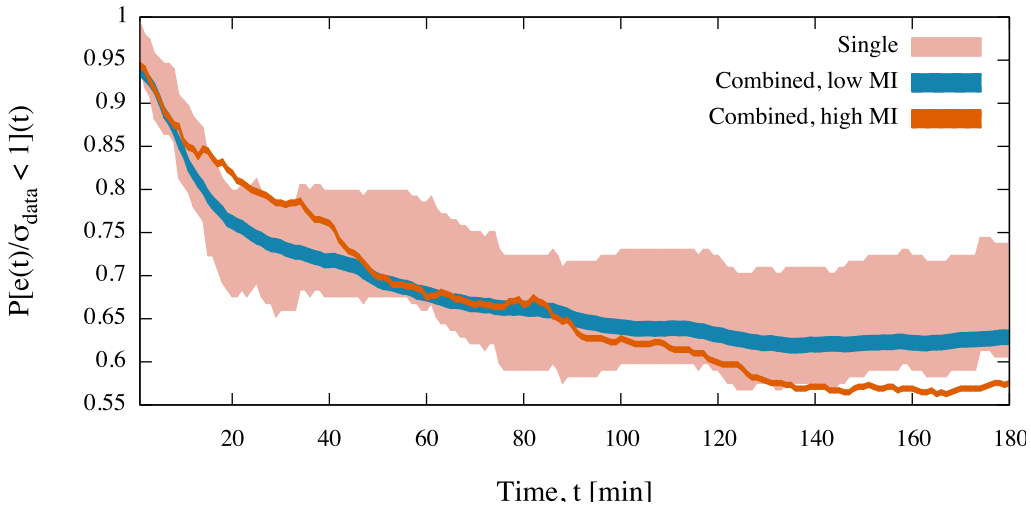


Figure 9: *Cabspotting dataset.* Temporal evolution of the fraction of mobility traces satisfying the condition $P(e(t)/\sigma_{data} < 1)$. Shaded areas indicate the 90% confidence bands around the average value.

## 5. Discussion

In the context of mobile applications, the prediction of mobility patterns of users is of great interest for several reasons. For instance, mobility forecasting could be used to determine where the person will be and who he/she will meet. Such an information can enable location-based mobile applications to provide personalised services relating to the context the user is in.

However, we are aware that there are scalability issues related to the implementation and the deployment of the proposed technique. In particular, it is well known that calculating mutual information in a multidimensional

environment (in this case, for a number of users larger than two) is computationally expensive and does not scale efficiently. In fact, in this case the computational complexity scales as $\mathcal{O}(N^n)$, where $N$ is the subset of users and $n$ is the cardinality of the tuple taken into account. However, the problem we are dealing with usually involves no more than 100 mobility traces (e.g., the size of the circle of most significant friends for an individual). For this reason, we can still evaluate mutual information values for any pair of traces, which scales as $\mathcal{O}(N^2)$. Nonetheless, the multivariate embedding reconstruction is not feasible for a phase-space larger than 40-dimensional. Even for a 2-coordinate signal representing a mobility trace, it is not unusual to have a large embedding reconstruction due to noisy data. Hence, no more than three mobility traces should be considered simultaneously.

It is worth noting that many factors could be considered as signals of social ties, according to the context of the deployment scenarios. As a consequence, the quality of predictions might be deeply affected, either positively or negatively, by the criteria used to detect social ties. In the Nokia MDC dataset, we had no information about social ties between individuals, neither of real nor virtual nature. In the available dataset, the presence of an individual in the address book of another one actually represents the strongest definition of a social tie. Moreover, two individuals with no social ties might show similar mobility patterns, resulting in a high value of mutual information. However, it is worth noticing again that the presence of a social tie is only a sufficient but not necessary condition for having a significant correlation between two mobility patterns. In fact, it could happen that two individuals are not socially linked (i.e., they are not friends, co-workers and so on), but their mobility is highly correlated, e.g., in the case of an individual whose work depends on the activities of another one and so on. If this is the case, the multivariate nonlinear predictor will greatly benefit of such a correlation, as in the case of the Cabspotting dataset. On the other hand, it could happen that two individuals with social ties show uncorrelated mobility patterns, degrading the accuracy of the predictor. It is likely that individuals with strong social ties (students, friends, co-workers and so on) behave similarly and their mobility traces are characterised by patterns with a high value of mutual information. Hence, the accuracy of the predictor will be improved in the case the dynamics of traces is highly correlated, even if a social tie does not exist.

A possible refinement of this work is the use of multivariate nonlinear prediction with non-uniform embedding (different delays) and local polynomial

fitting [30] in order to increase the accuracy of the prediction.

## 6. Conclusions

In this paper, we have shown discussed multivariate nonlinear time series techniques can be successfully applied to improve the prediction of movement of users, by considering the movement of people with correlated mobility patterns. More specifically, through the analysis of the Nokia Mobile Data Challenge traces, we have shown that it is possible to exploit the correlation of social interactions and user movement in order to improve the accuracy of forecasting of the future geographic position of a user. By means of the Cabspotting dataset we have also shown that when social ties information is not available, mutual information can be used to select pairs of users in order to improve prediction accuracy.

In other words, mobility correlation, measured by means of mutual information, and the presence of social ties can be used to improve movement forecasting by exploiting mobility data of other individuals. This correlation can be used as an indicator of potential existence of physical or distant social interactions and vice versa.

## References


[1] D. Lazer et al., Computational Social Science, Science 323 (2009) 721–723.

[2] D. Brockmann, L. Hufnagel, T. Geisel, The scaling laws of human travel, Nature 439 (2006) 462–465.

[3] S. Scellato, M. Musolesi, C. Mascolo, V. Latora, A. T. Campbell, NextPlace: A Spatio-temporal Prediction Framework for Pervasive Systems, in: Proceedings of Pervasive'11, San Francisco, California, USA.

[4] M. C. González, C. A. Hidalgo, A.-L. Barabási, Understanding individual human mobility patterns, Nature 453 (2008) 779–782.

[5] L. Aalto, N. Göthlin, J. Korhonen, T. Ojala, Bluetooth and WAP Push Based Location-aware Mobile Advertising System, in: Proceedings of MobiSys '04, pp. 49–58.

[6] T. M. Cover, J. A. Thomas, Elements of Information Theory, Wiley, 2006.

[7] M. De Domenico, A. Lima, M. Musolesi, Interdependence and Predictability of Human Mobility and Social Interactions, in: Proceedings of the Nokia Mobile Data Challenge 2012 Workshop, Newcastle, United Kingdom.

[8] J. K. Laurila, D. Gatica-Perez, I. Aad, J. Blom, O. Bornet, T.-M.-T. Do, O. Dousse, J. Eberle, M. Miettinen, The Mobile Data Challenge: Big Data for Mobile Computing Research, in: Proceedings of the Mobile Data Challenge by Nokia Workshop. Colocated with Pervasive'12.

[9] H. Kantz, T. Schreiber, Nonlinear time series analysis, Cambridge University Press, 1997.

[10] C. Chatfield, The Analysis of Time Series An Introduction, Chapman and Hall CRC, 2004.

[11] M. Piorkowski, N. Sarafijanovic-Djukic, M. Grossglauser, CRAWDAD trace set epfl/mobility/cab (v. 2009-02-24), Downloaded from http://crawdad.cs.dartmouth.edu/epfl/mobility/cab, 2009.

[12] F. Takens, Detecting strange attractors in turbulence, in: Dynamical Systems and Turbulence, Lecture Notes in Mathematics, volume 898, Springer-Verlag, 1981, pp. 366–381.

[13] I. Vlachos, D. Kugiumtzis, State space reconstruction from multiple time series, in: Topics on Chaotic Systems: Selected Papers from Chaos 2008 International Conference, World Scientific Publishing, p. 378.

[14] L. Cao, A. Mees, K. Judd, Dynamics from multivariate time series, Physica D 121 (1998) 75–88.

[15] H. Akaike, A New Look at the Statistical Model Identification, IEEE Transactions on Automatic Control C-19 (1974) 716–723.

[16] G. Schwarz, Estimating the Dimension of a Model, Ann. Stat. 6 (1978) 461–464.

[17] A. Fraser, H. Swinney, Independent coordinates for strange attractors from mutual information, Physical Review A 33 (1986) 1134.

[18] M. Kennel, R. Brown, H. Abarbanel, Determining embedding dimension for phase-space reconstruction using a geometrical construction, Physical Review A 45 (1992) 3403.

[19] R. Hegger, H. Kantz, Improved false nearest neighbor method to detect determinism in time series data, Physical Review E 60 (1999) 4970.

[20] S. Boccaletti, D. Valladares, L. Pecora, H. Geffert, T. Carroll, Reconstructing embedding spaces of coupled dynamical systems from multivariate data, Physical Review E 65 (2002) 35204–35204.

[21] R. Hegger, H. Kantz, T. Schreiber, Practical implementation of nonlinear time series methods: The TISEAN package, Chaos 9 (1999) 413–435.

[22] Farmer J.D. and Sidorowich J.J., Predicting chaotic time series, Phys. Rev. Lett. 59 (1987) 845–848.

[23] G. Sugihara, R. May, Nonlinear forecasting as a way of distinguish chaos from measurement error in time series, Nature 344 (1990) 734.

[24] M. Barahona, C. Poon, Detection of nonlinear dynamics in short, noisy time series, Nature 381 (1996) 215–217.

[25] T. Schreiber, A. Schmitz, Discrimination power of measures for nonlinearity in a time series, Phys. Rev. E 55 (1997) 5443–5447.

[26] M. De Domenico, V. Latora, Fast detection of nonlinearity and nonstationarity in short and noisy time series, EuroPhys. Lett. 91 (2010) 30005.

[27] Schreiber T. and Schmitz A., Improved surrogate data for nonlinearity tests, Phys. Rev. Lett. 77 (1996) 635.

[28] T. Schreiber, Constrained randomization of time series data, Phys. Rev. Lett. 80 (1998) 2105–2108.

[29] M. Casdagli, Nonlinear prediction of chaotic time series, Physica D 35 (1989) 335–356.

[30] L. Su, Prediction of multivariate chaotic time series with local polynomial fitting, Computers & Mathematics with Applications 59 (2010) 737–744.